

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE 1995		3. REPORT TYPE AND DATES COVERED Final Report
4. TITLE AND SUBTITLE Multiscale Methods in Molecular Mechanics			5. FUNDING NUMBERS F6170894W0939	
6. AUTHOR(S) Prof. Achi Brandt				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Department of Applied Mathematics and Computer Science, The Weizmann Institute of Science Rehovot Rehovot 76100 Israel			8. PERFORMING ORGANIZATION REPORT NUMBER N/A	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) EOARD PSC 802 BOX 14 FPO 09499-0200			10. SPONSORING/MONITORING AGENCY REPORT NUMBER SPC 94-4110	
11. SUPPLEMENTARY NOTES				
12a. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.			12b. DISTRIBUTION CODE A	
13. ABSTRACT (Maximum 200 words) This report results from a contract tasking Department of Applied Mathematics and Computer Science, The Weizmann Institute of Science as follows: Develop multi-scale methods using parallel processing capabilities for linking atomistic level simulations and continuum level material or structural modeling.				
14. SUBJECT TERMS Nil			15. NUMBER OF PAGES 23	
			16. PRICE CODE N/A	
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT UL	

Special Project SPC-94-4110
CONTRACTOR/PURCH ORDER NO. F6170894W0939

Multiscale Methods in Molecular Mechanics

Achi Brandt
Department of Applied Mathematics and Computer Science
The Weizmann Institute of Science
Rehovot, Israel 76100

First Quarterly Progress Report

Covering the period Nov. 1, 1994 to Jan. 31, 1995

DTIC QUALITY INSPECTED 5

Delivered to:

Dr. Ruth Pachter
Material Directorate
WL/MLPJ
Wright-Patterson AFB
OH 45433-6533

19980312 099

Multiscale Methods in Molecular Mechanics

A. Brandt

Department of Applied Mathematics and Computer Science
The Weizmann Institute of Science
Rehovot, 76100, Israel

1. Introduction

To maintain the lead in the area of developing large and complex highly specialized macromolecules with controlled properties, the fastest and most cost effective way is to design them first theoretically and then synthesize, process and characterize those candidates having sought after properties. However, current computational methods severely restrict the modeling effort since only relatively small molecular systems can be studied at atomic detail. The introduction of more advanced massively parallel supercomputers will not by itself radically change this situation, since the complexity of current algorithms rise very steeply with the problem size.

A sequence of preliminary model studies has indicated that the various types of scaling difficulties encountered in computational chemistry can in principle be overcome by multiscale algorithms. In collaboration with Dr. Ruth Pachter at the Materials Directorate, WL/MLPJ, Wright-Patterson AFB it has been decided to focus the present effort forward developing and demonstrating multiscale methods in molecular mechanics. Future, related efforts, depending on suitable support will include multiscale ab-initio calculations of electronic structures and the even more basic and accurate modeling of chemical reactions by multigrid Monte-Carlo methods for real-time Feynman path integrals.

It was decided that the fastest way to start the present project off is to use the contract money to hire Dr. Dov Bai (an American citizen) to work on it with Prof. Brandt, most of the time at the Weizmann Institute. Although Dr. Bai, with M.Sc. degree in physics and Ph.D. in applied mathematics, had little previous involvement with molecular mechanics, he has had extensive experience in multiscaling projects in several other scientific computation fields.

The actual work started on November 1, 1994. The three past months were devoted to the following.

(1) *Learning.* Dr. Bai has repeated computer experiments reported in [1] and practiced some of the previously-developed multiscale methods for fast force summations (cf. Sec. 3.1 below). Bai and Brandt have studied those parts of the molecular mechanics literature which seemed most relevant for the project: an annotated list is given in Sec. 2 below.

(2) *Outlining the long range objectives of the project,* based on the closer acquaintance with the subject obtained from our literature studies. See a summary of these objectives in Sec. 3 below.

(3) *Initial work on multiscale energy minimization methods.* Detailed approaches and development stages were planned. They are described in Sec. 4 below.

2. Studied Bibliography

A considerable amount of time was spent on getting acquainted with the field of molecular dynamics in general and dynamics of macromolecules in particular.

The classical works of L. Verlet [1] and A. Rahman [2] were useful in getting to know the basic capabilities of molecular dynamics methods in providing information about physical properties of systems in thermal equilibrium. Ref. [8] provided information about non-multileveled algorithms for summation of potentials and forces.

For larger molecules, Secs. II and V of [4] provided useful and comprehensive information about many types of molecular interaction terms as well as various techniques used in studying molecular statics and dynamics.

The publications of T. Schlick with coworkers (Refs. [6]–[10]) were very helpful in getting acquainted with the more recent attempts at increasing the time steps by using Newton and normal mode techniques to solve implicit schemes. These studies are most relevant to our current project as a reference point, being a similar attempt at fast algorithms for minimizing similar energy functionals.

We have also been delighted to discover several articles, such as [11]–[13], where some rudimentary forms of multiscaling have already been attempted, motivated by physical insights, with considerable gains. Unaware of several basic multiscale techniques, though, these attempts do not go far enough.

3. Long Term Objectives

Our study of the molecular mechanics literature has enabled us to outline in

detail the basic long term objectives of the project. They can be summarized as the following six tasks, ordered roughly according to the chronological order in which we have started or intend to continue to work on them, although much back and forth interaction between these coupled tasks is of course expected.

3.1 Fast summation of forces

Direct calculation of all the electrostatic interactions between N particles costs CN^2 computer operations, where C is around 10. Instead, several methods exist to sum the forces in just C_1N operations (see, e.g., survey [GR]), although note that in three dimensions $C_1 > 10^4$, so these methods become advantageous only for $N > 10^3$. A multiscale method for fast evaluation, suggested in [OS] (based on an idea described earlier in [G2] and [TO]), will be used by us. It is based on a decomposition of the two-particle potential into a local part and a smooth part, the latter being evaluated at larger scales (interpolated from coarser grids), where a similar decomposition is being recursively used. One possible advantage of this approach is considerably smaller values of C_1 . More important, the involved decompositions will give the kind of multiscale description of the force fields which is needed for the efficient multiscaling of the other tasks described below; see in particular the use of this decomposition in Secs. 4.2, 4.5 and 4.7 below.

Since we have already acquired enough preliminary experience with this force summation task, we will not develop it immediately further, until the need indeed arises in conjunction with our other tasks, described next.

3.2 Fast macromolecular energy minimization

This task serves two somewhat different objectives: one in statics, the other in dynamics. In statics, the objective is to calculate the lowest energy or the most stable conformations of large N -atom molecular structures; to find, in other words, the atom conformation $r = (r_1, r_2, \dots, r_N)$ for which the potential energy $E(r)$ is minimal. In dynamics, the objective is the solution of the system of equations arising at each time step of *implicit* dynamic simulations.

"Implicit" refers to the method which evaluates the forces (or the gradient of $E(r)$) at each time step (partly or wholly) in terms of the particle *arrival* positions, i.e., positions at the *end* of the step. This method ensures stability of very large time steps, but it does not yield the arrival positions explicitly. Instead, they should be calculated by solving a large system of equations. (Also, this method damp molecular vibrations at scales not resolved by the large time step; we return to this point below.) Solving the implicit system of equations is equivalent to minimizing an *augmented* energy functional, identical to $E(r)$ except for an additional quadratic *kinetic term* (cf., e.g., [6]). For large time steps this

additional term is locally very small, but its large-scale effect is still profound.

Starting from a close enough first approximation to the desired minimal energy configuration, current minimization methods would converge to that configuration, in the case of a strongly coupled N -atom structure, in $O(N^3)$ computer operations (not even including the operations for electrostatic force calculations, discussed above). Multiscale methods promise to obtain a solution in just $O(N)$ operations. This, however, requires a very careful development. Since this is our main *current* work, it is described in detail separately (Sec. 4 below). That description also includes a discussion of multiscale approaches to the harder situation where the given first approximation is *not* so close (Secs. 4.5–4.9).

3.3 Normal mode analysis

Near a given configuration of an N -atom molecular structure, the normal modes are the harmonic vibrations, i.e., the eigenfunctions of the harmonic (quadratic) approximation to the energy functional. They are important to approximate equilibrium and dynamic properties. In particular, the high-frequency modes may describe the vibrations not resolved by the large implicit time steps (see above), while low modes may describe important large-scale behavior.

The multiscale calculation of many low modes is expected to be extremely efficient, since it will be done on the *coarser levels* of the multiscale fast-solver described below (Sec. 4). Moreover, on those coarser levels one can also directly calculate the combined action of all these modes, and thus, due to the FAS version (see Sec. 4.6), calculate large-scale *anharmonic* effects as well. (This becomes rather similar to performing many large-scale statistical passes: see idem (d) in Sec. 3.4.)

We plan to demonstrate such basic capabilities after the multiscale solver (discussed in Sec. 4 below) is sufficiently developed.

In a less obvious way, a multiscale structure can also be very efficient in calculating higher modes, including the highest ones, since the separation between increasingly closer modes can be done on increasingly coarser levels. Simple models indicate that a well structured multiscale eigenbase for an N -atom molecule may be calculated in as little as $O(N \log N)$ operations. Such a calculation is, however, rather complicated, and its importance is questionable, because the exact large-scale separation of close *high-frequency* modes may have little physical significance. Probably more relevant to the high-frequency behavior are the statistical approaches, discussed next (Secs. 3.4 and 3.5).

3.4 Monte-Carlo methods at equilibrium

In equilibrium at temperature T , the probability of each configuration r to appear is proportional to $P(r) = \exp(-E(r)/T)$. To calculate equilibrium statistics, an atom-by-atom Monte-Carlo process is usually performed. In this process, each atom in its turn changes position stochastically, according to the probability density distribution $P(r)$. Making repeated sweeps of this process, one can calculate the desired statistics on the sequence of produced configuration.

To calculate accurate averages of some observable, however, an extremely long sequence of configurations is needed. There are two basic reasons for this complexity: (1) Due to the local nature of the Monte-Carlo process, only very slowly it affects large-scale conformational features, hence extremely many Monte-Carlo sweeps are needed to produce each new, statistically independent configuration. (2) Many such independent samples are needed to average out the deviation observed at each of them.

For some model problems, multigrid Monte-Carlo algorithms were developed which overcome *both* these complexity reasons (see [BG] and [L1]). The algorithms are similar to the multiscale cycles described below (Sec. 4), with the following four modifications.

(a) The energy functional should be the true potential $E(r)$, not the augmented one (cf. Sec. 3.2).

(b) The Gauss-Seidel relaxation (atom-by-atom minimization) sweeps should be replaced by atom-by-atom Monte-Carlo sweeps.

(c) The approximation of the Hamiltonian (energy functional) E in the coarsening process should be done in a stochastic manner, to retain the statistical fidelity (the "detailed balance"). Methods to achieve that are highly nontrivial, and may require careful research and development. However, in view of the approximate nature of the molecular-mechanics Hamiltonian to begin with, exact detailed balance may not be required, as long as statistical fidelity is retained in the limit of very smooth fluctuations. This can be achieved much more easily.

(d) The multiscale cycle should switch many times back and forth between coarse levels, before returning to finer levels. In this way many samples of large-scale features can be averaged over. Not so many passes are needed at the finer scales, because many fine-scale features are already present, and hence averaged over, in any one configuration.

3.5 Small-scale statistics with large-scale dynamics

The multiscale structure allows the combination of statistical simulations at small scales with time-accurate dynamics at large scales. For this purpose the

multiscale solver (Sec. 4) should be modified in two ways.

First, the time-step discretization should be such that it gives accurate (non-damping, energy conserving) approximations for all scales whose time-accurate dynamics need be simulated.

Secondly, at all finer scales (finer levels of the multiscale solver), the Gauss-Seidel relaxation sweeps should be replaced with Monte-Carlo sweeps, as described above (Sec. 3.4).

This scheme comes close to the one described, with a slightly different motivation, in Sec. 4.9 below.

3.6 Material "homogenization"

The coarse-level energy functional of a well-developed multiscale solvers should effectively yield the large-scale behavior of the simulated material. This would yield the crucial link between the atomistic level description and the continuum-level material modeling, thus enabling a basic understanding and design of phenomena at various scales.

Indeed, to simulate a material at macroscopic scales, a multiscale processing needs not resolve the atomic scales over a macroscopic domain. Rather, a tiny domain (but still very large compared with the smallest inter-atomic distances) can first be simulated. This simulation should employ 4D (space + time) periodic boundary conditions, using a 4D multiscaling; i.e., periodicity and coarsening is used not only in space, but also in time. (The optimal adjustment of the period size in each space direction and in time can very inexpensively be performed at the coarsest levels of such a solver.) In the next stage the periodicity is used to extend the domain (e.g., double the period size in each direction), while the description is *coarsened* (i.e., the finest scale of the multiscale solver is dropped, retaining the energy functional it supplied to the coarser levels). Then a simulation with this coarsened descriptions erase the original periodicity, retaining periodicity only at the extended domain boundaries. Repeating this procedure, the algorithm can simulate ever larger domains with ever coarser descriptions, until macroscopic scales are attained.

We hope to demonstrate simple instances of such multiscale expansions in some, much later stages of the project.

4. Energy Minimization

The first crucial step in constructing a fast multiscale energy minimizer is

to design its inner cycle: such that it yields fast convergence near the minimum, i.e., for sufficiently close first approximations. We discuss this cycle first (Secs. 4.1–4.4). Later (Secs. 4.5–4.9) we will discuss the nonlinear aspects of driving the cycles when the initial approximation is not so good.

When linearized, the molecular energy minimization problem is somewhat similar to the minimization problem encountered in structural mechanics, for which very efficient multigrid solvers have been developed. Of those, the closest to the ones needed in molecular mechanics are the *algebraic multigrid* (AMG) solvers [A1]–[A4], which do not assume that the problem arises from a continuum structure or that the unknowns are really placed on a grid. Some of the basic ideas of the AMG solvers, as well as the experience obtained with them, will be important in our present development.

4.1 Main difficulty and a model

There are several types of computational difficulties peculiar to molecular energy functionals. To develop an efficient multiscale solvers we must start by studying each of these difficulties in separation from the others, by treating appropriate model problems.

The main special difficulty arising in (linearized) molecular structures is that different kinds of forces are associated with very different coupling strengths (manifested for example in different time scales of the corresponding dynamics). This has profound implications for the multiscale cycle.

To explain and develop the principles of efficient cycles under such situations, we have started with a simple 2D model problem: A chain of N atoms at the planar positions $r_i = (x_i, y_i)$, ($i = 1, \dots, N$), with the energy functional

$$E(r) = \sum_{i=2}^N B_i (|r_i - r_{i-1}| - b_i)^2 + \sum_{i=2}^{N-1} K_i (\theta_i - \theta_{i0})^2,$$

where θ_i is the angle between the vector $r_{i+1} - r_i$ and the vector $r_i - r_{i-1}$, and where the bond constants B_i/b_i^2 are much larger than the angle constants K_i . To make the problem nontrivial, some of the atom positions (e.g., r_1 and r_N) may be fixed. This two-dimensional model mimics the three-dimensional situation where dihedral angle couplings are much weaker than those of both bonds and bond angles. (We will of course insist on applying to this model methods which do not depend on its special, non-representative features. For example, we will avoid using the bond lengths and angles as the dependent variables, which would yield a straightforward separation between weak and strong interactions, but would not be applicable in more general situations. Similarly we will refrain from solution techniques tailored only for one dimensional problems. For example, the technique of [11] can be very effective for such problems, but much less generally).

In the presence of such highly nonuniform couplings, the simple atom-by-atom minimization procedure (equivalent to Gauss-Seidel relaxation) is extremely inefficient. For example, in the model problem, if the K_i 's are comparable to ϵ times the typical size of B_i/b_i^2 , with some $0 < \epsilon \ll 1$, then the convergence of an atom-by-atom relaxation process would require $O(N^3/\epsilon)$ computer operations. Since coupling strengths in molecules range over at least four orders of magnitude, $\epsilon \approx 10^{-4}$ can be considered typical.

For small ϵ , the naive multigrid solvers (using the usual coarse-to-fine interpolation schemes) would also be quite ineffective, requiring $O(N/\epsilon)$ operations. By contrast, we believe that the general multigrid principles outlined below should solve any (harmonic) molecular energy minimization in just $O(N)$ operations, independently of coupling strength ratios. We describe below these general principles and, as an example, their specific application to the model problem.

4.2 Relaxation

The relaxation is a local process whose purpose in multiscale solvers is just the fast reduction of "non-smooth" errors, or their fast replacement by "smooth" ones. A given error function is considered "non-smooth" (or "locally reducible") if it is associated with relatively large residual forces, where "relatively large" means that they are comparable to the largest residual forces producible by any error comparable in magnitude to the given one.

The usual Gauss-Seidel (GS) relaxation (or atom-by-atom minimization) is actually very efficient in reducing non-smooth errors.

The only concern is the amount of work in calculating long-range (e.g., electrostatic forces). In the GS relaxation, the forces summed for each atom movement are based on the most recent locations of all other atoms. For long range forces, such a separate summation for each atom would be too expensive. On the other hand, the Jacobi-type relaxation, where the forces are based on atom positions at the beginning of the relaxation sweep and can therefore be based on efficient simultaneous summation of all forces, has smoothing properties much less satisfactory and requires carefully chosen under-relaxation parameters.

The best scheme seems indeed to be a combination of the Gauss-Seidel and Jacobi schemes, based on the decomposition of forces into a local part and a smooth part (see Sec. 3.1 above): The smooth part of the forces can be updated once per sweep, while the local forces should be updated during relaxation. Actually, since the local forces greatly diminish toward the margin of their support, it is suspected that only a fraction of them, those closest to the affected atom, need to be updated inside the relaxation sweep.

An interesting further possibility we want to explore is to entirely neglect the

smooth part of the forces during relaxation, since its significant effect is only on smooth errors. This would mean that the smooth part of the forces need not even be interpolated to the atom positions, implying a substantial work saving.

4.3 Coarse-level correction

The errors which are not reduced efficiently by the atom-by-atom minimization sweeps must exhibit relatively small residuals, as defined above. Hence such errors must be "smooth", or, more precisely, they nearly belong to the subspace spanned by the lower eigenfunctions (of the linearized system). They can thus be described by a smaller number of degrees of freedom. These degrees of freedom form the (first) "coarse level".

In the case of multigrid solvers for isotropic elliptic PDEs discretized on a uniform grid, relaxed errors (i.e., errors left after a couple of Gauss-Seidel relaxation sweeps) must be smooth in the usual sense (a grid function having small local differences compared with the function itself). Hence such errors can well be approximated by functions interpolated from a *coarser grid*. Thus the degrees of freedom of the "coarse level" are those of a coarser-grid functions. This function is calculated by solving the "coarse-grid equations", and its interpolation is then used as a correction to the former solution on the original (fine) grid. The coarse-grid equations themselves are determined by the requirement that the resulting interpolated correction lowers the energy as far as possible.

In the case of molecular structures, the sense of "smoothness" of relaxed errors is more complicated, due to the unevenness of the couplings and the lack of a uniform grid structure. Hence both the coarse level variables and the interpolation from them to form the corrections to the atom positions are also more complicated.

The coarse-level variables always represent displacements, i.e., changes, in atomic position (since they are designed to approximate the errors). One approach for choosing these variables is, as in AMG, to simply choose a *subset* of all the displacements of the original variables. We tend to take this approach for the first coarsening step discussed here. (By contrast, at coarser levels, where movements will no longer be constrained by very uneven couplings, the next-coarser level variables should probably best be defined on a 3D spatial *grid*, especially in the (usual) case where the molecules, including solvent, are not one dimensional, but fold and fill the full three dimensional space.)

The interpolation should directly reflect the strongest couplings. For example, in the above 2D model problem, the coarse level variables may simply include the displacement (vector) of every other atom, say those with even indices. The interpolation to the odd-numbered atoms is uniquely determined by the requirement that, to a first order, the bond lengths are unchanged.

In a more complex situation, the most general approach is to find the "sense of smoothness" of the relaxed errors by a local computation of the lowest eigenmodes. By this we mean, for example, computing the lowest eigenmodes of a system which only includes the interactions between the atoms inside a certain local box. In some interior subdomain of that box, a good coarse-to-fine interpolation can then be determined by the requirement that it is fully compatible with (exactly satisfied by) each of the lowest eigenmodes.

Although local in principle, this general procedure for deriving the interpolation is still quite expensive. Fortunately, at the finest (atomistic) level, a more explicit choice, based (as in the example) on the strongest couplings, will usually be quite straightforward and effective. (When the coarsening process is later applied at a coarser base level, with more complex interactions, it may be important to employ the general procedure, but its cost then will be much lower, due to the smaller number of variables at that coarser base level.)

Once the coarse-to-fine interpolation has been chosen, the *coarse-level energy functional* follows directly from the given (fine-level) energy functional. More precisely, a polynomial approximation (a Taylor expansion) for the latter (around the current fine grid solution) yields a polynomial (of the same degree) approximation for the former, since interpolation is a linear operator.

After minimizing this coarse-level energy functional (by the method discussed in the next section), the minimizing function is interpolated to the fine-level and used as a correction to the previous solution there. Then the process may be repeated: some Gauss-Seidel relaxation sweeps are made, followed by a new calculation of the coarse-level energy functional, its minimization, and interpolation of the minimum to further correct the atomic positions. And so on, producing fast converging iterations.

4.4 Recursion: multilevel cycle

The method for fast approximate minimization of the *coarse-level* energy is the same as that we have described for the original (atomistic) level. Namely, a couple of Gauss-Seidel relaxation sweeps (now on the coarse level) are followed by calculation of an energy functional for the next coarser level, (approximate) minimization of the latter, interpolation of the obtained minimum as a correction to the (first) coarse-level configuration, and perhaps some additional Gauss-Seidel sweeps for the latter. In this way we can recursively define the *multilevel* (or "*multiscale*") cycle.

The main remaining task in constructing the cycle is to choose, at each *current* level, the next-coarser-level variables and the coarser-to-current interpolation operator. These choices will determine the efficiency of the algorithm. Some of the

main approaches for making them were indicated above. The development and testing of versatile and efficient coarsening schemes will be our major task during at least the first research year.

4.5 Nonlinear aspects

So far we have dealt with fast convergence to the minimum of polynomial (e.g., quadratic) approximations to the energy functional. To be sure, the atom-by-atom relaxation process is actually performed with the *full* energy functional (local-Newton Gauss-Seidel relaxation); but the coarsening process is done in terms of a polynomial (Taylor) expansion of the energy functional around some current configuration. The expansion does not really assume small displacements: In terms of *local* forces (including the local part of the electrostatic force, in the sense of the decomposition mentioned in Sec. 3.1), the expansion only assumes small *strains*, i.e., small displacement *differences*, or, in other words, *smooth* displacements. Since what we want the coarse level to perform are just such smooth moves, the expansion holds true even for large displacements (except for the smooth part of the electrostatic forces — discussed below).

However, when the initial configuration is not close enough to the desired minimum of the full energy functional, just making a couple of iterations of polynomial approximations (e.g., a couple of global Newton steps) may not work. We plan several kinds of multiscale devices to deal with the nonlinear, more global convergence; they are described in the next sections (Secs. 4.6–4.9).

4.6 FAS cycles

One device is the FAS (“Full Approximation Scheme”) version of multigrid algorithms. In this scheme, the coarse-level variables are shifted: instead of using the *displacement* variables (those which, upon interpolation, describe atom displacements), each FAS variable is the *sum* of a displacement variable and the corresponding atom location just before the coarsening process. Since that location has been assumed (in the polynomial energy expansion) to be fixed throughout the solution processes for the coarse level (including the processes on still coarser levels), this shift is just a simple additive constant shift for each variable. (Note that the FAS variable actually describes the full new location of the atom; hence the term “full approximation” scheme. But the interpolation to atoms not represented on the coarse grid is still done in terms of the *displacements*, which are calculated by back-shifting the FAS variables.) The advantage of the FAS variables is that they allow to write the coarse-level equations in a special way, such that the approximation holds for a much wider range of coarse-level displacements (see detailed description, for the case of PDEs, in Chapter 8 of [G2]). In essence,

the FAS allows taking into account nonlinear (anharmonic) interactions at the coarse levels, since it retains knowledge of the full approximation, not just its displacements.

This in particular applies to the smooth part of the electrostatic forces: In term of the FAS variables that part can retain its form in the coarse level. Then later, on that coarse level, this smooth part will again be decomposed into a local part and a smooth part, "local" and "smooth" now in a larger-scale sense; and in the next coarsening, from the first coarse level to the next coarser one, it will be that *new* (larger-scale) smooth part that will retain its form. And so on, to increasingly coarser levels with increasingly smoother parts of the electrostatic interaction. This will make the energy Taylor expansions (cf. Sec. 4.5). valid for large (smooth) displacements even in terms of the electrostatic forces. (This also reinforces the importance of the smooth + local decomposition of electrostatic forces, as discussed in Sec. 3.1 above.)

4.7 FAS-FMG algorithm

The actual multiscale solver that will be used at each time step of molecular dynamics is the *FAS-FMG solver* (or the "*F* cycle"; described, e.g., in [PR]). In it, an (approximate) energy functional is first transferred to increasingly coarser levels, in terms of FAS variables. Then a so-called FMG (= "full multigrid") solver is applied. This solver starts at the coarsest level and sequentially proceeds to increasingly finer "base levels". At each base level, a first approximation is obtained by interpolating the displacements from the previous (next coarser) base level. Then the approximation is improved by several multilevel cycles, of the type described in Sec. 4.4, except that they start at the current base level, not at the original (finest, atomistic) level. Then the solution is interpolated to the next, finer base level; and so on, until the finest level (of atoms) is reached and multilevel cycles are performed for it.

Actually, the experience with PDEs is that only *one* multilevel cycle at each base level is usually enough, since the initial approximation, obtained from the next coarser level, is already very good. One multilevel cycle is usually also enough at the finest level, for the same reason, making this solver very inexpensive. Effectively, the FAS-FMG algorithm makes global-Newton iterations unnecessary.

4.8 Continuations

A general powerful technique to deal with difficult nonlinear problems is to combine the FMG process with a *continuation* (*parameter embedding*) process (see Sec. 8.3.2 in [G2]). This means that at the initial stages of the FMG solver (at multilevel cycles for coarse base levels) some problem parameters are changed, to

allow large scale movements not to be excessively affected by small scale details. As the FMG process advances to even finer base levels, the problem parameters are gradually restored, to finally reach their true value when the algorithm reaches the finest (the target, atomistic) level.

In particular, a natural FMG-continuation process for molecular mechanics is to ignore, at each base level, the local part of the electrostatic interactions of the next finer base level. That is, in the formulation of the base level energy functionals *in the beginning of the FAS-FMG algorithm*, this part is omitted; it is however *not* omitted in any coarsening within any of the multilevel *cycles*.

A related technique is to use a continuation process with a sequence of *several* multilevel cycles for the *same* base level, the finest base level in particular. Natural continuation parameters in this respect are the stiffness coefficients of the strongest couplings. For example, in the above 2D model problem (Sec. 4.1), the values of B_i may be softened (lowered) when the approximate energy functional of the coarse level is calculated (see end of Sec. 4.3; this softening is especially important in case higher-than-quadratic Taylor expansions are used). As convergence is approached at subsequent cycles, the softening can be gradually taken out. The full-strength (unsoftened) coefficients should still be used in all the *relaxation* sweeps between coarsening steps. The softening at each coarsening step is such that it allows the expected size of the large-scale displacements not to be stifled by the strong couplings. (Although the effect of the latter vanishes to a first order if interpolation is designed as described in Sec. 4.3, the second-order effect may still be stifling if the coupling strength ratio is very large.) Since such a stifling (being a second order effect) disappears for sufficiently small displacements, the softening can indeed be gradually taken out as convergence is approached.

4.9 Stochasticity, combined with dynamics

A powerful method to escape local energy minima in search of a more global minimum is to add stochasticity to the minimization process. *Simulated annealing* is the most well known example. Since in molecular mechanics problems a multitude of local minima comes with multiscale attraction basins, a multiscale version, called *multilevel annealing* (see Sec. 4 of [RA]), can be much more effective (as shown in [RA] for spin glass problems).

The energy functional discussed thus far is the *augmented* one, including the quadratic *kinetic term* (see Sec. 3.2 above). This augmented energy would be used in the annealing processes at each time step. For the dynamics of large biomolecules, however, the most natural stochasticity, and possibly the only one to guarantee correct dynamics, is the one associated with the true temperature of the material. This suggests a multiscale process where on finer scales (those scales associated with vibrations not resolved by the time step), instead of the

Gauss-Seidel relaxation, a Monte-Carlo process (cf. Sec. 3.4) will be applied. It will use the relevant physical temperature and the bare (not augmented) potential energy. This process comes in fact close to that of Sec. 3.5, except that it may be simpler because it can gloss over the problem of detailed balance. Also, it may be desired to finish the process off by an annealing phase (letting the temperature decrease gradually to zero).

References

- [1] L. Verlet, Computer 'experiments' on classical fluids, I. Thermodynamical properties of Lennard-Jones molecules, *Phys. Rev.* **98** (1967).
- [2] A. Rahman, Correlations in the motion of atoms in liquid Argon, *Phys. Rev.* **136** (1964) A405.
- [3] M.P. Allen and D.J. Tildesley, *Computer Simulation of Liquids*, Oxford University Press, New York (1990).
- [4] B.R. Brooks, R.E. Bruccoleri, B.D. Olafson, D.J. States, S. Swaminathan and M. Karplus, CHARMM: A program for macromolecular energy, minimization and dynamics calculations, *J. Comput. Chem.* **4** (1982) 187.
- [5] W.F. van Gunsteren, H.J.C. Berendsen and J.A.C. Rullmann, Stochastic dynamics for molecules with constraints Brownian dynamics of *n*-alkanes, *Molecular Phys.* **44** (1981) 69.
- [6] C.S. Peskin and T. Schlick, Molecular dynamics by backward-Euler method, *Comm. Pure Appl. Math.* **42** (1989) 1001.
- [7] A.M. Nyberg and T. Schlick, Increasing the time step in molecular dynamics, *Chem. Phys. Lett.* **198** (1992) 538.
- [8] P. Derreumaux, G. Zhang, T. Schlick and B. Brooks, A truncated Newton minimizer adopted for CHARMM and biomolecular applications, *J. Comput. Chem.* **15** (1994) 532.
- [9] T. Schlick, S. Figueroa and M. Mezei, A molecular dynamics simulation of water droplet by the implicit Euler/Langevin scheme, *J. Chem. Phys.* **94** (1991) 2118.
- [10] G. Zhang and T. Schlick, The Langevin/implicit-Euler/normal-mode scheme for molecular dynamics of large time steps, *J. Chem. Phys.* **101** (1994) 4995.
- [11] M.H. Hao, S.C. Harvey, Analyzing the normal mode dynamics of macromolecules by the component synthesis method, *Biopolymers* **32** (1992), 1393.

- [12] M.E. Tuckerman, B.J. Berne and A. Rossi, Molecular dynamics algorithm for multiple time scales: systems with disparate masses, *J. Chem. Phys.* **94** (1991) 1465.
- [13] M.E. Tuckerman and B.J. Berne, Molecular dynamics in systems with multiple time scales: systems with stiff and soft degrees of freedom and with short and long range forces, *J. Chem. Phys.* **95** (1991) 8362.
- [A1] A. Brandt, S. McCormick and J. Ruge, Algebraic multigrid (AMG) for automatic multigrid solution with application to geodetic computations, Institute for Computational Studies, POB 1852, Fort Collins, Colorado, 1982.
- [A2] A. Brandt, S. McCormick and J. Ruge, Algebraic multigrid (AMG) for sparse matrix equations, in: *Sparsity and its Applications* (D.J. Evans, ed.), Cambridge University Press, Cambridge, 1984, pp. 257-284.
- [A3] A. Brandt, Algebraic multigrid theory: The symmetric case, *Appl. Math. Comp.* **19** (1986) 23-56.
- [A4] J. Ruge and K. Stüben, Algebraic multigrid, in: *Multigrid Methods* (S.F. McCormick, ed.), SIAM, Philadelphia, 1987, pp. 73-130.
- [BG] A. Brandt, M. Galun and D. Ron, Optimal multigrid algorithms for calculating thermodynamic limits, *J. Stat. Phys.* **74** (1994) 313-348.
- [G2] A. Brandt, Guide to multigrid development, in: *Multigrid Methods* (W. Hackbusch and U. Trottenberg, eds.), Springer-Verlag, 1982, pp. 220-312.
- [GR] L. Greengard, Fast algorithms for classical physics, *Science* **265** (1994) 909-914.
- [L1] A. Brandt, Multigrid methods in lattice field computations, *Nucl. Phys. B* (Proc. Suppl.) **26** (1992) 137-180.
- [OS] A. Brandt, Multilevel computations of integral transforms and particle interactions with oscillatory kernels, *Comp. Phys. Comm.* **65** (1991) 24-38.
- [PR] A. Brandt and J. Greenwald, Parabolic multigrid revisited, in: *Multigrid Methods III* (W. Hackbusch and U. Trottenberg, eds.), Birkhäuser Verlag, Basel, 1991, pp. 143-154.
- [RA] A. Brandt, D. Ron and D.J. Amit, Multi-level approaches to discrete-state and stochastic problems, in: *Multigrid Methods II* (W. Hackbusch and U. Trottenberg, eds.), Springer Verlag, 1986, pp. 66-99.
- [TO] A. Brandt and A.A. Lubrecht, Multilevel matrix multiplication and fast solution of integral equations, *J. Comp. Phys.* **90** (1990) 348-370.

Special Project SPC-94-4110
CONTRACTOR/PURCH ORDER NO.\ F6170894W0939
Multiscale Methods in Molecular
Mechanics
Achi Brandt
Department of Applied Mathematics and Computer
Science
The Weizmann Institute of Rehovot, Israel 76100

Second Quarterly Progress Report
Covering the period Feb.\ 1, 1995 to Apr.\ 30, 1995
Delivered to: Dr. Ruth Pachter
Material Directorate WL/MLPJ
Wright-Patterson AFB OH 45433--6533

General

The former Report [1] has outlined the general research strategy we have developed for introducing fast multiscale methods into the computation of molecular dynamics, statics and statistics. The first crucial step in this strategy is the design of a multiscale cycle for fast energy minimization in the neighborhood of the minimum, i.e., when starting with sufficiently close first approximation. This cycle will serve to solve efficiently the large system of equations arising at each time step of dynamics simulations when an implicit time discretization is used (allowing very large time steps). Together with other techniques outlined in [1], including real-temperature stochastic steps at fine levels, this cycle will also be pivotal in energy minimizations starting $\{\textit{it farV}\}$ from the minimum. The coarsening techniques developed for such a cycle will later also be central for other molecular dynamics tasks sketched in [1], including fast normal mode analysis, multiscale Monte-Carlo processes at equilibrium, combination of small-scale statistics with large-scale dynamics, and derivation of bulk material properties.

Latest work

The main work over the last several months involved a sequence of technical steps aimed at a systematic development and examination of some of the basic numerical devices required for multiscaling of large molecular structures. Mostly we worked on designing $\{\textit{it coarsening principles}\}$, the core of any multiscale scheme.

A coarsening step is a procedure for transforming the system from a given fine level to the next coarser one. The objective of the step is to create a system (the coarser level) with substantially reduced (e.g., halved) number of degrees of freedom, but in which the same large-scale motions can still be described, associated with energy differences similar to those in the given (the fine) system. The construction of each such coarsening step consists of three inter-connected aspects: the choice of the coarse-level degrees of freedom, the design of the coarse-to-fine interpolation, and the derivation (and simplification) of the coarse-level energy functional. Each of these aspects depends on the

special features of the given fine level.

The dominant special feature of the {\it finest} level at hand (the given molecular structure, with all its degrees of freedom) is that different kinds of forces are associated with very different coupling strengths. To develop a general approach for dealing with such a situation, various simplified models exhibiting this feature were investigated (see for example the model in Sec.\ 4.1 of [1]). For each model, to separately study each of the above three aspects of coarsening, we have concentrated on two-level tests with ``unigrid" interpolations and/or ``unigrid" energy calculation.

In {\it two-level} tests, only one coarsening step at a time is studied (minimizing the coarse-level energy exactly, disregarding the question of how to do this efficiently). In ``{\it unigrid}" experiments, the fine-level changes associated with each pointwise coarse-level motion are immediately introduced along with that motion (instead of introducing all the fine-level changes after completing the coarse-level energy minimization). In ``{\it unigrid" interpolations}, the fine level changes associated with a coarse-level motion are defined by a (possibly long) process of local energy minimization (instead of some simple chosen interpolation). In ``{\it unigrid" energy calculations}, the energy difference associated with each coarse-level motion is calculated by explicitly finding the energy difference of the corresponding fine-level changes. All these types of simplified experiments involve of course procedures much less efficient than the ultimate algorithm, but they help to investigate the three aspects of coarsening separately from each other.

Tentative conclusions

Although our investigations are not yet completed, some preliminary conclusions already emerge. The most important one is that the first coarsening step (the only one we thoroughly researched) should be very efficient indeed: the two-level tests show very fast convergence (more than an order-of-magnitude residual-force reduction per cycle). This efficiency critically depends, however, on the correct choice of the interpolation.

We have developed a general computational approach for deriving the interpolation weights from the relative displacements resulting in local sequences of dummy relaxation steps. This approach is not restricted to the simple models we have used to develop and test it.

Another conclusion is that the nature of couplings can completely change from level to level. For example, they can become much more uniform and isotropic than in the base level.

Next task

Our next task will be to extend our studies to increasingly coarser levels and more realistic models. We will do it by introducing further features one by one, repeating for each of them the type of artificial cycles described above, as well as other technical steps, designed to separate out the development of coarsening principles into a systematic investigation of one aspect at a time.

References

[1] A. Brandt: Multiscale methods in molecular mechanics.

First Quarterly Progress Report, Air Force Special Project SPC-94-4110.

Special Project SPC-94-4110
CONTRACTOR/PURCH ORDER NO. F6170894W0939

Multiscale Methods in Molecular Mechanics

Achi Brandt
Department of Applied Mathematics and Computer Science
The Weizmann Institute of Science
Rehovot, Israel 76100

Third and Fourth Quarterly Progress Report

Covering the period May 1, 1995 to October 31, 1995

Progress Reports
are acceptable
please authorize
3rd, 4th payment.
J. Stein
22 Nov 95

Delivered to:

Dr. Ruth Pachter
Material Directorate
WL/MLPJ
Wright-Patterson AFB
OH 45433-6533

Third and Fourth Progress Report

Multiscale Methods in Molecular Mechanics

A. Brandt

Department of Applied Mathematics and Computer Science
The Weizmann Institute of Science
Rehovot, 76100, Israel

1. Review

The former Reports have outlined the general research strategy we have developed for introducing fast multiscale methods into the computation of molecular dynamics, statics and statistics (see [1] and a brief summary in Sec. 1 of [2]). The main work so far has involved a sequence of technical steps aimed at a systematic development and examination of *relaxation principles* (see Sec. 2 below) and *coarsening principles* (Sec. 3), together constituting the core of any multiscale scheme. Also a new *stochastic dynamics* approach has been advanced, which will allow very large time steps with natural thermalization of all the unresolved high-frequency modes (Sec. 4). A workshop for teaching and discussing these investigations and related research has been held at the Weizmann Institute (see Sec. 5).

2. Relaxation Principles

2.1 Relaxation as a solver

A theoretical study of a simple model of an N -atom polymer shows that minimizing the energy through simple atom-by-atom relaxation steps requires $O(N^3\epsilon^{-2})$ computer operations, where $\epsilon \sim 10^{-2}$ is the ratio between short and long time scales associated with, respectively, the strongest (bond length) and weakest (torsion) local couplings. We have shown that this amount of work can be reduced to $O(N^3)$ by relaxing simultaneously μ atoms at a time. The number μ is such that μ atoms could simultaneously still move even in the limit case of rigid bond lengths and rigid angles. For simple chains, at last $\mu = 4$ is required to obtain ϵ -independent convergence, and increasing μ up to $\mu = 6$ (allowing the simultaneous move to be fully three dimensional) would still substantially improve the convergence rate. Increasing μ beyond 6 would be a waste of effort.

2.2 Relaxation as a smoother

As a smoother in a multiscale solver, however, relaxation does *not* require such simultaneous moves. (Moreover, of course, the multiscale target complexity is $O(N)$, not $O(N^3)$.) On the other hand the smoother can be made substantially more efficient by being done in terms of *internal coordinates*. Thus, for example, in a simple chain, for each relaxed atom at its turn the energy to be minimized should be written (and the forces be linearized) in terms of the atom distances from its two neighbors and its angle of rotation around the axis through them. Only such steps efficiently smooth not just the bond-length errors, but also the bond-angle errors.

3. Coarsening Principles

A coarsening step is a procedure for transforming the system from a given fine level to the next coarser one. The objective of the step is to create a system (the coarser level) with substantially reduced (e.g., halved) number of degrees of freedom, but in which the same large-scale motions can still be described, associated with energy differences similar to those in the given (the fine) system. The construction of each such coarsening step consists of three inter-connected aspects: the choice of the coarse-level degrees of freedom, the design of the coarse-to-fine interpolation, and the derivation (and simplification) of the coarse-level energy functional. Each of these aspects depends on the special features of the given fine level.

The dominant special feature of the *finest* level at hand (the given molecular structure, with all its degrees of freedom) is that different kinds of forces are associated with very different coupling strengths. To develop a general approach for dealing with such a situation, various simplified models exhibiting this feature were investigated. For each model, to separately study each of the above three aspects of coarsening, we have focussed our research on two-level tests with “unigrid” interpolations and/or “unigrid” energy calculation (see explanations in [2]). Some of the main findings are the following.

For typical molecular structures we have found that efficient coarsening can be obtained by taking as the coarse-level degrees of freedom a certain *subset of the atomic positions*. In particular, for simple chains, one can simply choose every other atom to serve as a coarse atom. In the usual case of length, angle and torsion couplings, an even more efficient choice is to include in the coarse level the first *two* atoms from each subsequent disjoint *quintuple*. With this choice, for any given coarse-level positions, the location of all atoms is uniquely determined by the stronger couplings alone (i.e., by the bond-length and bond-angle couplings only).

The general computational approach for deriving the coarse-to-fine interpolation which was reported in [2] has been replaced by a new one, which is easier and

faster to apply. Briefly, the scheme is as follows. The three-vector *displacement* (i.e., position *change*) of each atom is interpolated from the coarse-level-atom displacements. The three-by-three-matrix interpolation coefficients are uniquely determined by the relations prevailing when a *localized* energy functional (the functional obtained by including only interactions between atoms in a certain neighborhood) is minimized.

In summary, we have developed most of the tools necessary for efficient coarsening procedures. For simple models of long chains with typical length, angle and torsion bond interactions, very fast convergence (close to an order-of-magnitude residual-force reduction per multigrid cycle) has been obtained.

4. Stochastic Dynamics

The need to execute very large time steps δt raises the question of how to model those oscillatory molecular modes whose period is not resolved by δt . We have developed a new approach, called *stochastic dynamics*, which “thermalizes” such modes in a natural way, particularly compatible with the multiscale framework.

Our starting point is a usual (deterministic) implicit-time-step discretization of Newton dynamics, which generates a set of equations that should be solved at each time step. This set of equations is equivalent to the minimization of a functional $H(x)$, having the general form

$$H(x) = E(x) + T(x, x_0, v_0),$$

where x and x_0 are the vectors of atomic positions at the end and at the beginning of the time step, respectively, v and v_0 are the corresponding velocity vectors, the relation between velocities and positions being given by

$$\frac{x - x_0}{\delta t} = v + \alpha(v_0 - v), \quad 0 \leq \alpha \leq \frac{1}{2},$$

$E(x)$ is the potential energy, and T is a certain “kinetic term”, quadratic in x . Such a time step would damp down all the unresolved molecular vibrations, thereby severely distorting also the large-scale dynamics, which depend on those vibrations for its heat-bath supply of energy and stochasticity. Instead, the new approach is to replace the minimization of $H(x)$ by a Monte-Carlo choice of x , with the Boltzmann-like probability distribution

$$P(x) = \frac{1}{Z} e^{-\beta H(x)},$$

where $\beta = (k_B T)^{-1}$, k_B is the Boltzmann constant, T is the absolute temperature, and Z is a normalizing constant. This Monte-Carlo choice of x can be

performed very efficiently by a multiscale cycle using methods analogous to those being developed for the deterministic minimization.

It can be shown that oscillatory modes with periods small compared with δt are fully thermalized (slip into equilibrium statistics) within one such stochastic-dynamics time step, thereby acting as the required heat bath, while large scale modes still assume nearly deterministic Newtonian evolution. Also, the multiscale stochastic simulation is expected to be easier than the deterministic energy minimization for very large time steps, at which the minimization process would often be plagued with a multitude of false local minima. The multiscale stochastic simulation employing the natural temperature of the material is likely to stride much more efficiently over that landscape of many local minima embedded in multiscale cascades of attraction basins.

We have started a detailed study of this approach, including normal-mode analyses and numerical simulations for simple models.

5. Workshop

Together with Prof. Tamar Schlick from Courant Institute of Mathematics at New York University, we have organized a workshop entitled, "Multigrid Tutorial, with Applications to Molecular Dynamics", held on October 10-12, 1995 here at the Weizmann Institute. The research supported by the present EOARD contract has been reported in detail by Brandt and Bai to an audience of about 35 investigators from the US, Israel, Germany and France, in addition to basic teaching of the multigrid background and several lectures on other computational approaches to molecular mechanics, delivered by others. See [3] for details.

References

- [1] A. Brandt: Multiscale methods in molecular mechanics. First Quarterly Progress Report, Air Force Special Project SPC-94-4110. Appeared also as Gauss Center Report WI/GC-3, April, 1995.
- [2] A. Brandt: Multiscale methods in molecular dynamics. Second Quarterly Progress Report, Air Force Special Project SPC-94-4110.
- [3] Proceedings of the workshop, "Multigrid Tutorial, with Applications to Molecular Dynamics", The Gauss Minerva Center for Scientific Computation at the Weizmann Institute of Science, Rehovot, Israel, October 10-12, 1995.